

**CLAIMS:**

1 A host coupled to a cluster fabric including one or more fabric-attached I/O  
2 controllers, comprising:  
3 a processor;  
4 a memory coupled to the processor; and  
5 an operating system provided with an I/O bus abstraction for the cluster fabric to report  
6 multiple paths to a target fabric-attached I/O controller.

1 2. The host as claimed in claim 1, wherein said operating system further comprises:  
2 a kernel; and  
3 a fabric bus driver to provide said I/O bus abstraction to the kernel for the cluster fabric to  
4 report the multiple paths to the target fabric-attached I/O controller.

1 3. The host as claimed in claim 2, wherein said fabric bus driver presents the cluster  
2 fabric to the kernel as a local I/O bus, and presents one or more target fabric-attached I/O  
3 controllers to the kernel as local I/O controllers.

1 4. The host as claimed in claim 2, further comprising a host-fabric adapter provided  
2 to interface the host to the cluster fabric.

1           5.     The host as claimed in claim 4, further comprising a fabric adapter device driver  
2 provided to control operation of the host-fabric adapter.

1           6.     The host as claimed in claim 5, wherein said fabric bus driver creates a separate  
2 device object for each port of the host-fabric adapter that can be used to communicate with the  
3 target fabric-attached I/O controller and establish the multiple paths to the target fabric-attached  
4 I/O controller.

1           7.     The host as claimed in claim 5, wherein said multiple paths are utilized for load  
2 balancing I/O requests and/or for fault tolerance when one or more paths to the target fabric-  
3 attached I/O controller fail.

1           8     The host as claimed in claim 5, wherein said fabric bus driver creates a single  
2 device object for the target fabric-attached I/O controller even if multiple ports of the host-fabric  
3 adapter can be used to communicate with the target fabric-attached I/O controller.

1           9.     An operating system for a host coupled to a cluster fabric including one or more  
2 fabric attached I/O controllers, comprising:  
3 a kernel;

1 an I/O manager operatively coupled to the kernel;  
2 one or more I/O controller drivers operatively coupled to the kernel, each controller driver  
3 specific for a specific type of I/O controller; and  
4 a fabric bus driver operatively coupled to the I/O manager to provide an I/O bus  
5 abstraction to the I/O manager for the cluster fabric to report multiple paths to a target fabric-  
6 attached I/O controller.

10. The operating system as claimed in claim 9, wherein said fabric bus driver appears  
to the I/O manager as a local I/O bus driver.

11. The operating system as claimed in claim 9, wherein said fabric driver presents the  
cluster fabric to the I/O manager as a local I/O bus and presents the one or more fabric attached  
I/O controllers as local I/O controllers connected to the local I/O bus.

12. The operating system as claimed in claim 9, further comprising one or more local  
I/O bus drivers operatively coupled to the I/O manager.

13. The operating system as claimed in claim 12, wherein said local I/O bus drivers and  
said fabric bus driver communicate with the I/O manager using a common set of procedures.

1 14. A cluster comprising:  
2 a cluster fabric;  
3 a host including an operating system coupled to the cluster fabric;  
4 an I/O controller attached to the cluster fabric; and  
5 a fabric manager coupled to the cluster fabric, for assigning I/O controllers in the cluster  
6 fabric to at least said host and sending messages to said host indicating that the I/O controller has  
7 been assigned;

8 wherein said operating system including a fabric bus driver provided to report multiple  
9 paths to a target fabric-attached I/O controller.

10 15. The cluster as claimed in claim 14, wherein said operating system further  
11 comprises a kernel, and said fabric bus driver provided said I/O bus abstraction to the kernel for  
12 the cluster fabric to report the multiple paths to the target fabric-attached I/O controller.

13 16. The cluster as claimed in claim 14, further comprising a host-fabric adapter  
14 provided to interface the host to the cluster fabric, and a fabric adapter device driver provided to  
15 control operation of the host-fabric adapter.

16 17. The cluster as claimed in claim 15, wherein said fabric bus driver creates a separate  
17 device object for each port of the host-fabric adapter that can be used to communicate with the

1 target fabric-attached I/O controller and establish the multiple paths to the target fabric-attached  
2 I/O controller.

1 18. The cluster as claimed in claim 17, wherein said multiple paths are utilized for  
2 loading balancing I/O requests and/or for fault tolerance when one or more paths to the target  
3 fabric-attached I/O controller fail.

1 19. The cluster as claimed in claim 15, wherein said fabric bus driver creates a single  
2 device object for the target fabric-attached I/O controller even if multiple ports of the host-fabric  
3 adapter can be used to communicate with the target fabric-attached I/O controller.

4 20. The cluster as claimed in claim 14, wherein said fabric manager comprises:  
5 a fabric services to detect the connection or presence of the target fabric-attached I/O  
6 controller and to assign a network address to the target fabric-attached I/O controller; and  
an I/O controller manager coupled to the fabric services to assign the target fabric-  
attached I/O controller to said host and to send messages to said host indicating that the target  
fabric-attached I/O controller has been assigned.

1 21. A computer usable medium having computer readable program code means  
2 embodied therein for use in a host system to report multiple paths to a target fabric-attached agent

1 via a cluster fabric, said computer readable program code means comprising:

2 a fabric bus driver provided to create and report multiple paths to a target fabric-attached  
3 *ai* I/O controller via the cluster fabric, and  
*Cont*

4 a fabric adapter device driver provided to interface to the cluster fabric for enabling  
5 reporting the multiple paths to the target fabric-attached I/O controller.

*2017*  
1 22. A method of initializing a host to report multiple paths to a target agent via a  
2 cluster fabric, comprising:

3 loading an operating system kernel into a memory;

4 loading an I/O manager into the memory;

5 loading a local I/O bus driver and a fabric bus driver providing a local I/O bus abstraction  
6 for the cluster fabric into the memory;

7 enabling the local I/O bus driver to identify any local I/O controllers connected to a  
8 corresponding local I/O bus;

9 enabling the fabric bus driver to identify any fabric-attached I/O controllers assigned to the  
10 host, and report the identified local I/O controllers connected to the local I/O bus and the  
11 identified fabric-attached I/O controllers to the I/O manager;

12 loading an I/O controller driver into the memory for each reported I/O controller; and

13 enabling the fabric bus driver to create and report multiple paths to a target fabric-attached  
14 I/O controller via the cluster fabric.

1           23.     The method as claimed in claim 22, wherein said identified local I/O controllers  
2 connected to the local I/O bus and said identified fabric-attached I/O controllers to the I/O  
3 manager are reported using a common set of procedures or commands.

1           24.     A method of initializing a host to report multiple paths to a target agent via a  
2 cluster fabric, comprising:

3                   loading an operating system kernel into a memory;  
4                   loading an I/O manager into the memory; and  
5                   loading a local I/O bus driver and a fabric bus driver providing a local I/O bus abstraction  
6 for the cluster fabric into the memory;  
7                   enabling the local I/O bus driver to identify any local I/O controllers connected to a  
8 corresponding local I/O bus; and  
9                   enabling the fabric bus driver to identify any fabric-attached I/O controllers assigned to the  
10 host, identifies all paths to a target fabric-attached I/O controller, create one instance of an I/O  
11 controller driver stack for each path to the target fabric-attached I/O controller, and report all  
12 multiple paths to a target fabric-attached I/O controller via the cluster fabric.